REMARKS

Claims 1-55 are pending in the application. Claims 1, 16, 27, 42, 53, and 55 are independent.

Claims 1, 12, 23, 27, 38, 49, and 53 and their dependents stand rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. The Examiner has identified antecedent basis issues in these claims and they have been amended accordingly.

Claims 1-5, 10, 16, 21, 27, 36, 42, 47, and 53-55 stand rejected under 35 U.S.C. §103(a) as obvious over Allan et al. in view of Davies et al.

With regard to claim 1, the Examiner states that Allan et al. discloses a method for providing flow control (col. 3, lines 51-54, provide an addressing convention for carriage of ATM over Ethernet) for multiple signal streams (fig. 2, flows of cells 39, 39', 39") over a single ETHERNET link (fig. 2, E-Mux 21; col. 7, lines 5-8, E-Mux exchanges ATM cells with an ATM UNI, which is connected in turn to an ATM network... E-Mux 21 is also connected to an Ethernet LAN), comprising: receiving PDUs (protocol data units) (col. 7, line 35, cells) from multiple streams (col. 7, lines 35-39, flow of cells addressed to an end station 39, 39', 39") at a first MAC (media access control) client (fig. 2, E-Mux 21); encapsulating each PDU in a MAC frame (fig. 3A,

MAC frame 3) which includes an identification of the stream to which the PDU belongs (fig. 3A, VPINCI 40); transmitting the MAC frames over an ETHERNET link (fig. 2, arrows between E-Mux 21 and Ethernet network 25) to a second MAC client (fig. 2, frame manager 33); receiving the MAC frames at the second MAC client (col. 8, lines 30-40, frame manager 33 receives the frame from interface 47); decapsulating each PDU (col. 8, lines 30-40, provides the payload to PDU manager 31); forwarding each PDU (col. 8, lines 30-40, provides the payload to PDU manager 31).

Before turning to the Davies reference, it should be noted that the portions of Allan et al. cited by the Examiner do not teach what the Examiner says they teach. Col. 3 of Allan et al. does not teach or suggest any kind of "flow control" as that term is used in the art of telecommunications and certainly not as that term is used in the Applicant's specification. Allan et al. teaches an addressing scheme. While packets cannot "flow" without an address, "flow control" involves more than an addressing scheme. The term "flow control" means controlling the rate of data transmission.

The Examiner further states that Allan et al. teaches multiple signal streams over a single Ethernet link. The Examiner identifies the signal streams as 39, 39', etc. and the link as E-Mux 21 in Fig. 2 of Allan et al. The streams are identified by Allan et al. as "end station[s] in [an] Ethernet network 25". The E-Mux 21 is identified by Allan et al. as transferring traffic from the LAN 25 to the ATM network 23 and vice versa. It is not "a single Ethernet link" as claimed in claim 1. An Ethernet link has an Ethernet MAC client at both ends. Traffic enters the Ethernet link as an Ethernet frame and exits the Ethernet

link as an Ethernet frame in both directions. The E-Mux 21 described by Allan et al. receives ATM cells from the network 23 in the incoming direction and receives Ethernet frames from the network 25 in the outgoing direction. Thus, the E-Mux 21 cannot be considered an Ethernet link as that term is normally understood and as it is used in the Applicant's specification and claims.

The Examiner further states that Allan et al. discloses a first MAC client and identifies the first MAC client as being the E-Mux 21. This can't possibly be correct since the E-Mux 21 was previously identified by the Examiner as the single Ethernet link and it cannot be both. In fact, it is neither a MAC client nor a single Ethernet link.

Later in the Examiner's explanation of Allen et al., the Examiner identifies the single Ethernet link as "the arrows between [the] E-Mux 21 and [the] Ethernet network 25". Of course there are no arrows between 21 and 25 in Fig. 2 of Allan et al. The Applicant can only guess whether the Examiner is referring to the line between 47 and 21 or the line between 47 and 25.

The Examiner states that the second MAC client in claim 1 is shown by Allan et al. as the "frame manager 33 [which] receives the frame from interface 47". This interpretation is problematic because the Examiner previously identified first MAC client of claim 1 as being the E-Mux 21. Thus, under this interpretation, the second MAC client is a component of the first MAC client which is clearly an incorrect interpretation.

Reviewing claim 1 as addressed so far by the Examiner, the first MAC client receives PDUs from multiple streams, puts each into a MAC frame, transmits the MAC frames over an Ethernet link to the second MAC client. The Examiner has identified the end stations ES of the LAN 25 in Fig. 2 of Allan et al. as being the multiple streams of PDUs and the first MAC client as 21 in Fig. 2. After 21 receives the PDUs from network 25, those PDUs must be put into MAC frames in order to satisfy the limitations of claim 1, but that is not at all what happens in Fig. 2 of Allan et al. All of the PDUs received by 21 from 25 are taken out of MAC frames and put into ATM cells. Therefore, the Examiner's rejection must fail at this point without even considering the Davies et al. reference.

Even if one considers the flow of data in the other direction, Fig. 2 of Allan et al. cannot be made to read on the portions of claim 1 addressed so far. If one were to consider ATM cells received at 27 to be multiple signal streams as claimed in claim 1, the frame manager 33 (or the E-Mux 21) to be the first MAC client, and each of the end stations (ES) to be second MAC clients, the limitations of claim 1 still could not be met. Claim 1 requires multiple signal streams to be carried from a first MAC client over a single Ethernet link to a second MAC client where the multiple streams are decapsulated. In Fig. 2 of Allan et al. each separate stream of ATM cells travels inside an Ethernet frame to a separate destination ES via separate links. Different streams do not arrive at the same MAC client. There is also a major conceptual difference between Allan et al. and the present invention. In the present invention, non-Ethernet signals (e.g. SONET signals) arrive at a first location where they are encapsulated into Ethernet frames, the

frames are transported to a second location where the non-Ethernet signals are decapsulated from the frames and sent on to final destinations as non-Ethernet signals (e.g. SONET signals). In Allan et al. signals begin and end their journey as Ethernet frames. In Allan et al. Ethernet frames are encapsulated in ATM cells at one location, transported via ATM to a second location where they are decapsulated back into Ethernet frames.

Davies et al. teaches flow control and quality of service provisions for Frame Relay Protocols. Prioritized frames are allocated to transmission queues responsive to their priorities. Each queue has an associated subsidiary Ethernet MAC which transmits frames from its queue subject to a scheduler which selects from the set of MACs according to a pre-determined algorithm. The multiple logical paths between corresponding pairs of transmitter and receiver subsidiary MACs are preferably multiplexed over a single physical channel. If congestion occurs at the receiver, then Ethernet PAUSE frames may be sent back to the transmitter, directed to specific subsidiary MACs--typically those with lower priority--to suspend transmission from the corresponding queue for a time period indicated in the PAUSE frame. In this way, back-pressure flow control may be applied selectively so that large amounts of low priority traffic do not cause unnecessary delays to higher priority traffic.

While that description taken from the Abstract of Davies et al. may seem similar to the operation of PAUSE frames in the present invention, there are several important differences. First, Davies et al. buffers MAC frames in a queue prior to transmission.

Claim 1 specifies buffering PDUs which have been extracted from Frames at the end of an Ethernet link.

Second, in Davies et al., frames are transmitted from their queues according to a scheduler which schedules their transmission through a switch fabric according to an algorithm. In the present invention, all PDUs are treated equally.

Third, in Davies et al., the de-multiplexed frames are directed to output queues according to frame priority. In the present invention, PDUs are forwarded to port buffers associated with the stream identified in the MAC frame from which the PDU was decapsulated. In Davies et al., the multiplexer replaces MAC addresses with a single port MAC address "so that the whole of the input data appears to have issued from a single Ethernet input port." See ¶0048 of Davies et al.

Fourth, in Davies et al., the PAUSE frame is used to suppress transmission of traffic from a specific queue. That is, the PAUSE frame is sent to a specific MAC associated with a particular receiver queue and indicates a time period during which the transmitter is to suspend transmission. See ¶0052 of Davies et al. In the present invention, the PAUSE control frame indicates the fullness condition of each buffer.

For the foregoing reasons it is submitted that claim 1 and its dependents are allowable over Allan et al. in view of Davies et al.

Independent claim 16 requires receiving MAC frames from a MAC client, each frame containing a PDU and an indication of the stream to which the PDU belongs; decapsulating the PDUs and storing each in a buffer associated with the stream indicated in the frame; monitoring the fullness of each buffer; and transmitting a PAUSE control frame to the MAC client indicating the fullness condition of each buffer.

As argued above, Allan et al. does not teach or suggest any buffers. The buffers in Davies et al. contain MAC frames, not decapsulated PDUs. Furthermore, the PAUSE frames in Davies et al. do not indicate the fullness condition of a plurality of buffers. Thus, it should be clear that claim 16 and its dependents are allowable over Allan et al. in view of Davies et al.

Independent claim 27 is an apparatus claim which mirrors method claim 1 with means plus function language. Therefore, the arguments made above with regard to claim 1 apply as well to claim 27. Thus, it should be clear that claim 27 and its dependents are allowable over Allan et al. in view of Davies et al.

Independent claim 42 is an apparatus claim which mirrors method claim 16 with means plus function language. Therefore, the arguments made above with regard to claim 16 apply as well to claim 42. Thus, it should be clear that claim 42 and its dependents are allowable over Allan et al. in view of Davies et al.

Independent claim 53 is an apparatus claim which is slightly narrower than claim 27. Therefore, the arguments made above with regard to claims 1 and 27 apply as well to claim 53. Thus, it should be clear that claim 53 and its dependents are allowable over Allan et al. in view of Davies et al.

Independent claim 55 is an apparatus claim which is slightly narrower than claim 42. Therefore, the arguments made above with regard to claims 16 and 42 apply as well to claim 55. Thus, it should be clear that claim 55 is allowable over Allan et al. in view of Davies et al.

Claims 6, 7, 17, 18, 32, 33, 43, and 44 stand rejected under 35 U.S.C. §103(a) as obvious over Allan et al. in view of Davies et al. as applied to claim 1 and further in view of Bitner. These claims are all dependent claims and, as argued above, the combination of Allan et al. in view of Davies et al. does not obviate the claims from which these claims depend. Bitner does not disclose the elements of claim 1 missing from Allan et al. and Davies et al. Thus, this rejection is moot.

Claims 8, 9, 19, 20, 34, 35, 45, and 46 stand rejected under 35 U.S.C. §103(a) as obvious over Allan et al. in view of Davies et al. as applied to claim 1 and further in view of Chuah. These claims are all dependent claims and, as argued above, the combination of Allan et al. in view of Davies et al. does not obviate the claims from which these claims depend. Chuah does not disclose the elements of claim 1 missing from Allan et al. and Davies et al. Thus, this rejection is moot.

Claims 11-13, 22-24, 37-39, and 48-50 stand rejected under 35 U.S.C. §103(a) as obvious over Allan et al. in view of Davies et al. as applied to claim 1 and further in view of Ramakrishnan. These claims are all dependent claims and, as argued above, the combination of Allan et al. in view of Davies et al. does not obviate the claims from which these claims depend. Ramakrishnan does not disclose the elements of claim 1 missing from Allan et al. and Davies et al. Thus, this rejection is moot.

Claims 14, 15, 25, 26, 40, 41, 51, and 52 stand rejected under 35 U.S.C. §103(a) as obvious over Allan et al. in view of Davies et al. as applied to claim 1 and further in view of Ramakrishnan as applied to claim 13 and further in view of Erjanne. These claims are all dependent claims and, as argued above, the combination of Allan et al. in view of Davies et al. does not obviate the claims from which these claims depend. Erjanne does not disclose the elements of claim 1 missing from Allan et al., Davies et al., and Ramakrishnan. Thus, this rejection is moot.

In light of all of the above, it is submitted that the claims are in order for allowance, and prompt allowance is earnestly requested. Should any issues remain outstanding, the Examiner is invited to call the undersigned attorney of record so that the case may proceed expeditiously to allowance.

Respectfully submitted,

/David P. Gordon/

David P. Gordon
Reg. No. 29,996
Attorney for Applicant(s)

GORDON & JACOBSON, P.C.
60 Long Ridge Road
Suite 407
Stamford, CT 06902
voice: (203) 323-1800
fax: (203) 323-1803

November 27, 2007